

Convex Hulls under Uncertainty

Pankaj K. Agarwal*

Sariel Har-Peled[†]

Subhash Suri[‡]

Hakan Yildiz[§]

Wuzhou Zhang[¶]

June 26, 2014

Abstract

We study the convex-hull problem in a probabilistic setting, motivated by the need to handle data uncertainty inherent in many applications, including sensor databases, location-based services and computer vision. In our framework, the uncertainty of each input site is described by a probability distribution over a finite number of possible locations including a *null* location to account for non-existence of the point. Our results include both exact and approximation algorithms for computing the probability of a query point lying inside the convex hull of the input, time-space tradeoffs for the membership queries, a connection between Tukey depth and membership queries, as well as a new notion of β -hull that may be a useful representation of uncertain hulls.

1. Introduction

The convex hull of a set of points is a fundamental structure in mathematics and computational geometry, with wide-ranging applications in computer graphics, image processing, pattern recognition, robotics, combinatorics, and statistics. Worst-case optimal as well as output-sensitive algorithms are known for computing the convex hull; see the survey [Sei04] for an overview of known results.

In many applications, such as sensor databases, location-based services or computer vision, the location and sometimes even the existence of the data is uncertain, but statistical information can be used as a probability distribution guide for data. This raises the natural computational question: what is a robust and useful convex hull representation for such an uncertain input, and how well can we compute it? We explore this problem under two simple models in which both the location and the existence (presence) of each point is described probabilistically, and study basic questions such as what is the probability of a query point lying inside the convex hull, or what does the probability distribution of the convex hull over the space look like.

*Department of Computer Science; Duke University; Durham, NC, 27708, USA; pankaj@cs.duke.edu <<http://www.cs.duke.edu/>>; <http://www.cs.duke.edu/~protect/unhbox/voidb@x\penalty@M\{pankaj/>.

[†]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@uiuc.edu; <http://sarielhp.org/>.

[‡]Department of Computer Science; University of California; Santa Barbara; CA, 93106, USA; suri@cs.ucsb.edu; <http://www.cs.ucsb.edu/~protect/unhbox/voidb@x\penalty@M\{suri>.

[§]Department of Computer Science; University of California; Santa Barbara; CA, 93106, USA; hakan@cs.ucsb.edu; <http://www.cs.ucsb.edu/~protect/unhbox/voidb@x\penalty@M\{hakan>.

[¶]Department of Computer Science; Duke University; Durham, NC, 27708, USA; wuzhou@cs.duke.edu <<http://www.cs.duke.edu/>>; <http://www.cs.duke.edu/~protect/unhbox/voidb@x\penalty@M\{wuzhou/>.

Uncertainty models. We focus on two models of uncertainty: unipoint and multipoint. In the **unipoint model**, each input point has a fixed location but it only exists probabilistically. Specifically, the input \mathcal{P} is a set of pairs $\{(p_1, \gamma_1), \dots, (p_n, \gamma_n)\}$ where each p_i is a point in \mathbb{R}^d and each γ_i is a real number in the range $(0, 1]$ denoting the probability of p_i 's existence. The existence probabilities of different points are independent; $P = \{p_1, \dots, p_n\}$ denotes the set of sites in \mathcal{P} .

In the **multipoint model**, each point probabilistically exists at one of multiple possible sites. Specifically, \mathcal{P} is a set of pairs $\{(P_1, \Gamma_1), \dots, (P_m, \Gamma_m)\}$ where each P_i is a set of n_i points and each Γ_i is a set of n_i real values in the range $(0, 1]$. The set $P_i = \{p_i^1, \dots, p_i^{n_i}\}$ describes the possible sites for the i th point of \mathcal{P} and the set $\Gamma_i = \{\gamma_i^1, \dots, \gamma_i^{n_i}\}$ describes the associated probability distribution. The probabilities γ_i^j correspond to disjoint events and therefore sum to at most 1. By allowing the sum to be less than one, this model also accounts for the possibility of the point not existing (i.e. the **null** location)—thus, the multipoint model generalizes the unipoint model. In the multipoint model, $P = \bigcup_{i=1}^m P_i$ refers to the set of all sites and $n = |P|$.

Our results. The main results of our paper can be summarized as follows.

- (A) We show (in Section 2) that the membership probability of a query point $q \in \mathbb{R}^d$, namely, the probability of q being inside the convex hull of \mathcal{P} , can be computed in $O(n \log n)$ time for $d = 2$. For $d \geq 3$, assuming the input and the query point are in general position, the membership probability can be computed in $O(n^d)$ time. The results hold for both unipoint and multipoint models.
- (B) Next we describe two algorithms (in Section 3) to preprocess \mathcal{P} into a data structure so that for a query point its membership probability in \mathcal{P} can be answered quickly. The first algorithm constructs a **probability map** $\mathbb{M}(\mathcal{P})$, a partition of \mathbb{R}^d into convex cells, so that all points in a single cell have the same membership probability. We show that $\mathbb{M}(\mathcal{P})$ has size $\Theta(n^{d^2})$, and for $d = 2$ it can be computed in optimal $O(n^4)$ time. The second one is a sampling-based Monte Carlo algorithm for constructing a near-linear-size data structure that can approximate the membership probability with high likelihood in sublinear time for any fixed dimension.
- (C) We show (in Section 4) a connection between the membership probability and the Tukey depth, which can be used to approximate cells of high membership probabilities. For $d = 2$, this relationship also leads to an efficient data structure.
- (D) Finally, we introduce the notion of **β -hull** (in Section 5) as another approximate representation for uncertain convex hulls in the multipoint model: a convex set C is called **β -dense** for \mathcal{P} , for $\beta \in [0, 1]$, if C contains at least β fraction of each uncertain point. The β -hull of \mathcal{P} is the intersection of all β -dense sets for \mathcal{P} . We show that for $d = 2$, the β -hull of \mathcal{P} can be computed in $O(n \log^3 n)$ time.

Related work. There is extensive and ongoing research in the database community on uncertain data; see [DRS09] for a survey. In the computational geometry community, the early work relied on deterministic models for uncertainty (see e.g. [LÖ9]), but more recently probabilistic models of uncertainty, which are closer to the models used in statistics and machine learning, have been explored [ACTY09, AAH⁺13, Phi09, KCS11a, KCS11b, SVY13]. The convex-hull problem over uncertain data has received some attention very recently. Suri *et al.* [SVY13] showed that the problem of computing the most likely convex hull of a point set in the multipoint model is NP-hard. Even in the unipoint model, the problem is NP-hard for $d \geq 3$. They also presented an $O(n^3)$ -time algorithm for computing the most likely convex hull under the unipoint model in \mathbb{R}^2 . Zhao *et al.* [ZYN12] investigated the problem of computing the probability of each uncertain point lying on the convex hull, where they aimed to return the set of (uncertain) input points whose probabilities of being on the convex hull are at least

some threshold. Jørgensen *et al.* [JLP11] showed that the distribution of properties, such as areas or perimeters, of the convex hull of \mathcal{P} may have $\Omega(\prod_{i=1}^m n_i)$ complexity if all the sites lie on or near a circle.

2. Computing the Membership Probability

For simplicity, we describe our algorithms under the unipoint model, and then discuss their extension to the multipoint model. We begin with the 2D case.

2.1. The two-dimensional case

Let $\mathcal{P} = \{(p_1, \gamma_1), \dots, (p_n, \gamma_n)\}$ be a set of n uncertain points in \mathbb{R}^2 under the unipoint model. Recall that $P = \{p_1, \dots, p_n\}$ is the set of all sites of \mathcal{P} . For simplicity of description, we assume that the sites are in general position, i.e., no two share coordinates and no three are collinear. A subset $B \subseteq P$ is the outcome of a probabilistic experiment with probability

$$\gamma(B) = \prod_{p_i \in B} \gamma_i \times \prod_{p_i \notin B} \overline{\gamma}_i,$$

where $\overline{\gamma}_i$ is the complementary probability $1 - \gamma_i$. By definition, for a point q , the **probability** of q to lie in the convex-hull of B is

$$\mu(q) = \sum_{B \subseteq P \mid q \in \text{CH}(B)} \gamma(B),$$

where $\text{CH}(B)$ is the convex hull of B . This unfortunately involves an exponential number of terms. However, observe that for a subset $B \subseteq P$, the point q is **outside** $\text{CH}(B)$, if and only if q is a vertex of the convex hull $\text{CH}(B \cup \{q\})$. So, let $C = \text{CH}(B \cup \{q\})$, and V be the set of vertices of C . Then, we have that $\mu(q) = 1 - \Pr[q \in V]$.

If $B = \emptyset$, then clearly $C = \{q\}$ and $q \in V$. Otherwise, $|V| \geq 2$ and $q \in V$ implies that q is an endpoint of exactly two edges on the boundary of C .¹ In this case, the first edge following q in the counter-clockwise order of C is called the **witness edge** of q being in V . Hence, $q \in V$ if and only if $B = \emptyset$ or (exclusively) B has a witness edge, i.e.,

$$\Pr[q \in V] = \Pr[B = \emptyset] + \sum_{i=1}^n \Pr[qp_i \text{ is the witness edge of } q \notin \text{CH}(B)].$$

The first term can be computed in linear time. To compute the i th term in the summation, we observe that qp_i is the witness edge of B if and only if $p_i \in B$ and B contains no sites to the right of the oriented line spanned by the vector $\overrightarrow{qp_i}$, and the corresponding probability is $\gamma_i \cdot \prod_{p_j \in G_i} \overline{\gamma}_j$, where G_i is the set of sites to the right of $\overrightarrow{qp_i}$. This expression can be computed in $O(n)$ time. It follows that one can compute $1 - \mu(q)$, and therefore $\mu(q)$, in $O(n^2)$ time. The computation time can be improved to $O(n \log n)$ as described in the following paragraph.

Improving the running time. The main idea is to compute the witness edge probabilities in radial order around q . We sort all sites in counter-clockwise order around q . Without loss of generality,

¹If B consists of a single site p_i , then C is the line segment qp_i . In this case, we consider the boundary of C to be a cycle formed by two edges: one going from q to p_i , and one going from p_i back to q .




Figure 1: Sites in radial order around q .




Figure 2: The set W_i .

assume that the circular sequence p_1, \dots, p_n is the resulting order. (See Figure 1.) We first compute the probability that qp_1 is the witness edge in $O(n)$ time. Then, for increasing values of i from 2 to n , we compute the probability that qp_i is the witness edge by updating the probability for qp_{i-1} , in $O(1)$ amortized time. In particular, let W_i denote the set of sites in the open wedge bounded by the vectors $\overrightarrow{qp_{i-1}}$ and $\overrightarrow{qp_i}$. (See Figure 2.) Notice that $G_i = G_{i-1} \cup \{p_{i-1}\} \setminus W_i$. It follows that the probability for qp_i can be computed by multiplying the probability for qp_{i-1} with $\frac{\gamma_i}{\gamma_{i-1}} \times \frac{\gamma_{i-1}}{\prod_{p_j \in W_i} \gamma_j}$. The cost of a single update is $O(1)$ amortized because total number multiplications in all the updates is at most $4n$. (Each site affects at most 4 updates.) Finally, notice that we can easily keep track of the set W_i during our radial sweep, as changes to this set follow the same radial order.

Theorem 2.1. *Given a set of n uncertain points in \mathbb{R}^2 under the unipoint model, the membership probability of a query point q can be computed in $O(n \log n)$ time.*

2.2. The d -dimensional case

The difficulty in extending the above to higher dimensions is an appropriate generalization of witness edges, which allow us to implicitly sum over exponentially many outcomes without overcounting. Our algorithm requires that all sites, including the query point q , are in general position, i.e., no $k+1$ points of $\mathcal{P} \cup \{q\}$ lie on a $(k-1)$ -hyperplane when projected into a subset of k coordinates, where $2 \leq k \leq d$.


Let B be an outcome, $C = \text{CH}(B \cup \{q\})$ its convex hull, and V the vertices of C . Let $\lambda(B \cup \{q\})$ denote the point with the lowest x_d -coordinate in $B \cup \{q\}$. Clearly, if q is $\lambda(B \cup \{q\})$ then $q \in V$; otherwise, we condition the probability based on which point among B is $\lambda(B \cup \{q\})$. Therefore, we can write

$$\Pr[q \in V] = \Pr[q = \lambda(B \cup \{q\})] + \sum_{1 \leq i \leq n} \Pr[p_i = \lambda(B \cup \{q\}) \wedge q \in V].$$

It is easy to compute the first term. We show below how to compute each term of the summation in $O(n^{d-1})$ time, which gives the desired bound of $O(n^d)$.

Consider an outcome B with $p_i \in B$. Let B', p'_i and q' denote the projections of B, p_i and q respectively on the hyperplane $x_d = 0$, which we identify with \mathbb{R}^{d-1} . Let us define $C' = \text{CH}(B' \cup \{q'\}) \subset \mathbb{R}^{d-1}$, and let V' be the vertices of C' .

Let $\vec{r}(p'_i, q')$ denote the open ray emanating from q' in the direction of the vector $\overrightarrow{p'_i q'}$ (that is, this ray is moving “away” from p'_i). A facet f of C is a **p_i -escaping** facet for q , if q is a vertex of f and the projection of f on \mathbb{R}^{d-1} intersects $\vec{r}(p'_i, q')$. See the figure on the right. The following lemma is key to our algorithm. The points of C projected into $\partial C'$ form the ***silhouette*** of C .



Lemma 2.2. (A) If $q' \in V'$ then q is a silhouette vertex of C and vice versa.

(B) If $p_i \in B$ then q has at most one p_i -escaping facet on C .

(C) The point q is a non-silhouette vertex of the convex-hull C if and only if q has a (single) p_i -escaping facet on C .

Proof: (A) By definition.

(B) If q has a p_i -escaping facet then it is a vertex of the convex-hull C . Consider the union of facets adjacent to q , and observe that the projection of this “tent” can fold over itself in the projection only if q is on the silhouette. Specifically, if q is not on the silhouette then the claim immediately holds.

Otherwise, q is on the silhouette then the open ray $\vec{r}(p'_i, q')$ does not intersect C' , and there are no p_i -escaping facets.

(C) Follows immediately from (B), by observing that in this case, the projected “tent”, surrounds q' , and as such one of the facets must be an escaping facets for p_i . ■

Given a subset of sites $P_\alpha \subseteq P \setminus \{p_i\}$ of size $(d-1)$, define $f(P_\alpha)$ to be the $(d-1)$ -dimensional simplex $\text{CH}(P_\alpha \cup \{q\})$. Since $p_i = \lambda(B \cup \{q\})$ implies $p_i \in B$, we can use Lemma 2.2 to decompose the i th term as follows:

$$\begin{aligned} \Pr[p_i = \lambda(B \cup \{q\}) \wedge q \in V] &= \Pr[p_i = \lambda(B \cup \{q\}) \wedge q' \in V'] \\ &+ \sum_{\substack{P_\alpha \subseteq P \setminus \{p_i\} \\ |P_\alpha| = (d-1) \\ f(P_\alpha) \text{ is } p_i\text{-escaping for } q}} \Pr[p_i = \lambda(B \cup \{q\}) \wedge f(P_\alpha) \text{ is a facet of } C]. \end{aligned}$$

The first term is an instance of the same problem in $(d-1)$ dimensions (for the point q' and the projection of P), and thus is computed recursively. For the second term, we compute the probability that $f(P_\alpha)$ is a facet of C as follows. Let $G_1 \subseteq P$ be the subset of sites which are on the other side of the hyperplane supporting $f(P_\alpha)$ with respect to p_i . Let $G_2 \subseteq P$ be the subset of sites that are below p_i along the x_d -axis. Clearly, $f(P_\alpha)$ is a facet of C (and $p_i = \lambda(B \cup \{q\})$) if and only if all points in P_α and p_i exist in B , and all points in $G_1 \cup G_2$ are absent from B . The corresponding probability can be written as

$$\gamma_i \times \prod_{p_j \in P_\alpha} \gamma_j \times \prod_{p_j \in G_1 \cup G_2} \overline{\gamma_j}.$$

This formula is valid only if $P_\alpha \cap G_2 = \emptyset$ and p_i has a lower x_d -coordinate than q ; otherwise we set the probability to zero. This expression can be computed in linear time, and the whole summation term can be computed in $O(n^d)$ time. Then, by induction, the computation of the i th term takes $O(n^d)$ time. Notice that the base case of our induction requires computing the probability $\Pr[p_i = \lambda(B \cup \{q\}) \wedge q^{(d-2)} \in V^{(d-2)}]$ (where $^{(d-2)}$ indicates a projection to \mathbb{R}^2). Computing this probability is essentially a two-dimensional membership probability problem on q and P , but is conditioned on the existence of p_i and the non-existence of all sites below p_i along d th axis. Our two dimensional algorithm can be easily adapted to solve this variation in $O(n \log n)$ time as well. (Briefly, we apply the same algorithm but we ignore all points that are below p_i . We later adjust the Finally, we can improve the computation time for the i th term to $O(n^{d-1})$ by considering the facets $f(P_\alpha)$ in radial order. The details can be found in Appendix B.

Remark. The degeneracy of the input is easy to handle in two dimensions, but creates some technical difficulties in higher dimensions that we are currently investigating.

Theorem 2.3. *Let \mathcal{P} be an uncertain set of n points in the unipoint model in \mathbb{R}^d and q be a point. If the input sites and q are in general position, then one can compute the membership probability of q in $O(n^d)$ time, using linear space.*

Extension to the multipoint model. The algorithm extends to the multipoint model easily by modifying the computation of the probability for an edge or facet. Deferring the details to Appendix C, we conclude the following.

Theorem 2.4. *Given an uncertain set \mathcal{P} of n points in the multipoint model in \mathbb{R}^d and a point $q \in \mathbb{R}^d$, we can compute the membership probability of q in $O(n \log n)$ time for $d = 2$, and in $O(n^d)$ time for $d \geq 3$ if input sites and q are in general position.*

3. Membership Queries


We describe two algorithms – one deterministic and one Monte Carlo – for preprocessing a set of uncertain points for efficient membership-probability queries.

Probability map. The *probability map* $\mathbb{M}(\mathcal{P})$ is the subdivision of \mathbb{R}^d into maximal connected regions so that $\mu(q)$ is the same for all query points q in a region. The following lemma gives a tight bound on the size of $\mathbb{M}(\mathcal{P})$.

Lemma 3.1. *The worst-case complexity of the probability map of a set of uncertain points in \mathbb{R}^d is $\Theta(n^{d^2})$, under both the unipoint and the multipoint model, where n is the total number of sites in the input.*

Proof: We prove the result for the unipoint model, as the extension to the multipoint is straightforward. For the upper bound, consider the set H of $O(n^d)$ hyperplanes formed by all d -tuples of points in \mathcal{P} . In the arrangement $\mathcal{A}(H)$ formed by these planes, each (open) cell has the same value of $\mu(q)$. This arrangement, which is a refinement of $\mathbb{M}(\mathcal{P})$, has size $O((n^d)^d) = O(n^{d^2})$, establishing the upper bound.

For the lower bound, consider the problem in two dimensions; extension to higher dimensions is straightforward. We choose the sites to be the vertices p_1, \dots, p_n of a regular n -gon, where each site exists with probability γ , $0 < \gamma < 1$. See the figure on the right. Consider the arrangement \mathcal{A} formed by the line segments $p_i p_j$, $1 \leq i < j \leq n$, and treat each face as relatively open. If $\mu(f)$ denotes the membership probability for a face f of \mathcal{A} , then for any two faces f_1 and f_2 of \mathcal{A} , where f_1 bounds f_2 (i.e., $f_1 \subset \partial f_2$), we have $\mu(f_1) \geq \mu(f_2)$, and $\mu(f_1) > \mu(f_2)$ if $\gamma < 1$. Thus, the size of the arrangement \mathcal{A} is also a lower bound on the complexity of $\mathbb{M}(\mathcal{P})$. This proves that the worst-case complexity of $\mathbb{M}(\mathcal{P})$ in \mathbb{R}^d is $\Theta(n^{d^2})$. ■



We can preprocess this arrangement into a point-location data structure, giving us the following result for $d = 2$.

Theorem 3.2. *Let \mathcal{P} be a set of uncertain points in \mathbb{R}^2 , with a total of n sites. \mathcal{P} can be preprocessed in $O(n^4)$ time into a data structure of size $O(n^4)$ so that for any point $q \in \mathbb{R}^d$, $\mu(q)$ can be computed in $O(\log n)$ time.*

Appendix D describes how to construct the data structure in $O(n^4)$ time.

Remark. For $d \geq 3$, due to our general position assumption, we can compute the membership probability only for d -faces of $\mathbb{M}(\mathcal{P})$, and not for the lower-dimensional faces. In that case, by utilizing a point-location technique in [Cha93], one can build a structure that can report the membership probability of a query point (inside a d -face) in $O(\log n)$ time, with a preprocessing cost of $O(n^{d^2+d})$.

Monte Carlo algorithm. The size of the probability map may be prohibitive even for $d = 2$, so we describe a simple, space-efficient Monte Carlo approach for quickly approximating the membership probability, within absolute error. Fix a parameter $s > 1$, to be specified later. The preprocessing consists of s rounds, where the algorithm creates an outcome A_j of \mathcal{P} in each round j . Each A_j is preprocessed into a data structure so that for a query point $q \in \mathbb{R}^d$, we can determine whether $q \in \text{CH}(A_j)$.

For $d \leq 3$, we can build each $\text{CH}(A_j)$ explicitly and use linear-size point-location structures with $O(\log n)$ query time. This leads to total preprocessing time $O(sn \log n)$ and space $O(sn)$. For $d \geq 4$, We use the data structure in [MS92] for determining whether $q \in A_j$, for all $1 \leq j \leq s$. For a parameter t such that $n \leq t \leq n^{\lfloor d/2 \rfloor}$ and for any constant $\sigma > 0$, using $O(st^{1+\sigma})$ space and preprocessing, it can compute in $O(\frac{sn}{t^{1/\lfloor d/2 \rfloor}} \log^{2d+1} n)$ time whether $q \in \text{CH}(A_j)$ for every j .

Given a query point $q \in \mathbb{R}^d$, we check for membership in all $\text{CH}(A_j)$, and if it lies in k of them, we return $\hat{\mu}(q) = k/s$ as our estimate of $\mu(q)$. Thus, the query time is $O(\frac{sn}{t^{1/\lfloor d/2 \rfloor}} \log^{2d+1} n)$ for $d \geq 4$, $O(s \log n)$ for $d = 3$, and $O(\log n + s)$ for $d = 2$ (using fractional cascading).

It remains to determine the value of s so that $|\mu(q) - \hat{\mu}(q)| \leq \varepsilon$ for all queries q , with probability at least $1 - \delta$. For a fixed q and outcome A_j , let X_i be the random indicator variable, which is 1 if $q \in \text{CH}(A_j)$ and 0 otherwise. Since $\mathbb{E}[X_i] = \mu(q)$ and $X_i \in \{0, 1\}$, using a Chernoff-Hoeffding bound on $\hat{\mu}(q) = k/s = (1/s) \sum_i X_i$, we observe that $\Pr[|\hat{\mu}(q) - \mu(q)| \geq \varepsilon] \leq 2 \exp(-2\varepsilon^2 s) \leq \delta'$. By Lemma 3.1, we need to consider $O(n^{d^2})$ distinct queries. If we set $1/\delta' = O(n^{d^2}/\delta)$ and $s = O((1/\varepsilon^2) \log(n/\delta))$, we obtain the following theorem.

Theorem 3.3. *Let \mathcal{P} be a set of uncertain points in \mathbb{R}^d under the multipoint model with a total of n sites, and let $\varepsilon, \delta \in (0, 1)$ be parameters. For $d \geq 4$, \mathcal{P} can be preprocessed, for any constant $\sigma > 0$, in $O((t^{1+\sigma}/\varepsilon^2) \log \frac{n}{\delta})$ time, into a data structure of size $O((t^{1+\sigma}/\varepsilon^2) \log \frac{n}{\delta})$, so that with probability at least $1 - \delta$, for any query point $q \in \mathbb{R}^d$, $\hat{\mu}(q)$ satisfying $|\mu(q) - \hat{\mu}(q)| \leq \varepsilon$ and $\hat{\mu}(q) > 0$ can be returned in $O(\frac{n}{t^{1/\lfloor d/2 \rfloor} \varepsilon^2} \log \frac{n}{\delta} \log^{2d+1} n)$ time, where t is a parameter and $n \leq t \leq n^{\lfloor d/2 \rfloor}$. For $d \leq 3$, the preprocessing time and space are $O(\frac{n}{\varepsilon^2} \log \log \frac{n}{\delta} \log n)$ and $O(\frac{n}{\varepsilon^2} \log \frac{n}{\delta})$, respectively. The query time is $O(\frac{1}{\varepsilon^2} \log(\frac{n}{\delta}) \log n)$ (resp. $O(\frac{1}{\varepsilon^2} \log \frac{n}{\delta})$) for $d = 3$ (resp. $d = 2$).*

4. Tukey Depth and Convex Hull

The membership probability is neither a convex nor a continuous function, as suggested by the example in the proof of Lemma 3.1. In this section, we establish a helpful structural property of this function, intuitively showing that the probability stabilizes once we go deep enough into the “region”. Specifically, we show a connection between the Tukey depth of a point q with its membership probability; in two dimensions, this also results in an efficient data structure for approximating $\mu(q)$ quickly within a small absolute error.

Estimating $\mu(q)$. Let Q be a set of weighted points in \mathbb{R}^d . For a subset $A \subseteq Q$, let $w(A)$ be the total weight of points in A . Then the **Tukey depth** of a point $q \in \mathbb{R}^d$ with respect to Q , denoted by $\tau(q, Q)$,

is $\min w(Q \cap H)$ where the minimum is taken over all halfspaces H that contain q .² If Q is obvious from the context, we use $\tau(q)$ to denote $\tau(q, Q)$. Before bounding $\mu(q)$ in terms of $\tau(q, Q)$, we prove the following lemma.

Lemma 4.1. *Let Q be a finite set of points in \mathbb{R}^d . For any $p \in \mathbb{R}^d$, there is a set $\mathcal{S} = \{S_1, \dots, S_T\}$ of d -simplices formed by Q such that (i) each S_i contains p in its interior; (ii) no pair of them shares a vertex; and (iii) $T \geq \lceil \tau(p, Q)/d \rceil$.*

Proof: As long as $\tau(p, Q) > 0$, $p \in \text{CH}(Q)$, and by Carathéodory Theorem [Eck93], there is a d -simplex S with its $d+1$ vertices in Q such that $p \in S$. Remove the vertices of S from Q , and repeat the argument. Let S_1, \dots, S_T be the resulting simplices. Observe that at most d vertices of S can be in an halfspace passing through p , which implies that the Tukey depth of p drops by at most d after each iteration of this algorithm. Hence $T \geq \lceil \tau(p, Q)/d \rceil$. ■

We now use Lemma 4.1 to bound $\mu(p)$ in terms of $\tau(p, P)$.

Theorem 4.2. *Let \mathcal{P} be a set of n uncertain points in the uniform unipoint model, that is, each point is chosen with the same probability $\gamma > 0$. Let P be the set of sites in \mathcal{P} . There is a constant $c > 0$ such that for any point $p \in \mathbb{R}^d$ with $\tau(p, P) = t$, we have $(1 - \gamma)^t \leq 1 - \mu(p) \leq d \exp\left(-\frac{\gamma t}{cd^2}\right)$.*

Proof: For the first inequality, fix a closed halfspace H that contains t points of P . If none of these t points is chosen then p does not appear in the convex hull of the outcome, so $1 - \mu(p) \geq (1 - \gamma)^t$.

Next, let \mathcal{S} be the set of simplices of Lemma 4.1, and let V be its set of vertices, where $T \geq \lceil t/d \rceil$. Let $n' = |V| = (d+1)T$. Set $\varepsilon = \frac{1}{d+1}$. A random subset of V of size $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon\delta}) = O(d^2 \log \frac{d}{\delta})$ is an ε -net for halfspaces, with probability at least $1 - \delta$.

In particular, any halfspace passing through p , contains at least T points of V . That is, all these halfspaces are ε -heavy and would be stabbed by an ε -net. Now, if we pick each point of V with probability γ , it is not hard to argue that the resulting sample R is an ε -net³. Indeed, the expected size (and in with sufficiently large probability) of $R \cap V$ is $n'' = n'\gamma = (d+1)T\gamma \geq t\gamma$. As such, for some constant c , we need the minimal value of δ such that the inequality $t\gamma \geq cd^2 \ln \frac{d}{\delta}$ holds, which is equivalent to $\exp\left(\frac{t\gamma}{cd^2}\right) \geq \frac{d}{\delta}$. This in turn is equivalent to $\delta \geq d \exp\left(-\frac{t\gamma}{cd^2}\right)$. Thus, we set $\delta = d \exp\left(-\frac{t\gamma}{cd^2}\right)$.

Now, with probability at least $1 - \delta$, for a point p in \mathbb{R}^d , with Tukey depth at least t , we have that p is in the convex-hull of the sample. ■

Remark. Theorem 4.2 can be extended to the multipoint model. Assuming that each uncertain point has n_i sites and each site is chosen with probability γ , one can show that $(1 - \gamma)^t \leq 1 - \mu(p) \leq d \exp\left(-\frac{\gamma t}{cd^2 n^*}\right)$, where $n^* = \max_{1 \leq i \leq m} n_i$.

Theorem 4.2 can be extended to the case when each point p_i of \mathcal{P} is chosen with different probability, say, γ_i . In order to apply Theorem 4.2, we convert \mathcal{P} to a multiset \mathcal{Q} , as follows. We choose a parameter $\eta = \frac{\delta}{10n}$. For each point $p_i \in \mathcal{P}$, we make $w_i = \left\lceil \frac{\ln(1 - \gamma_i)}{\ln(1 - \eta)} \right\rceil$ copies of p_i , each of which is selected with probability η . We can apply Theorem 4.2 to \mathcal{Q} and show that if $\tau(q, \mathcal{Q}) \geq \frac{d^2}{\eta} \ln(2d/\delta)$, then $\mu(q, \mathcal{Q}) \geq (1 - \delta/2)$. Omitting the further details, we conclude the following.


²If the points in Q are unweighted, then $\tau(q, Q)$ is simply the minimum number of points that lie in a closed halfspace that contains q .

³The standard argument uses slightly different sampling, but this is a minor technicality, and it is not hard to prove the ε -net theorem with this modified sampling model.

Corollary 4.3. Let $\mathcal{P} = \{(p_1, \gamma_1), \dots, (p_n, \gamma_n)\}$ be a set of n uncertain points in \mathbb{R}^d under the unipoint model. For $1 \leq i \leq n$, set $w_i = \left\lceil \frac{\ln(1-\gamma_i)}{\ln(1-\delta/10n)} \right\rceil$ be the weight of point p_i . If the (weighted) Tukey depth of a point $q \in \mathbb{R}^d$ in $\{p_1, \dots, p_n\}$ is at least $\frac{10d^2n}{\delta} \ln(2d/\delta)$, then $\mu(q, \mathcal{P}) \geq 1 - \delta$.

Data structure. Let \mathcal{P} be a set of points in the uniform unipoint model in \mathbb{R}^2 , i.e., each point appears with probability γ . We now describe a data structure to estimate $\mu(q)$ for a query point $q \in \mathbb{R}^2$, within additive error $1/n$. We fix a parameter $t_0 = \frac{c}{\gamma} \ln n$ for some constant $c > 0$. Let $\mathcal{T} = \{x \in \mathbb{R}^2 \mid \tau(x, \mathcal{P}) \geq t_0\}$ be the set of all points whose Tukey depth in P is at least t_0 . \mathcal{T} is a convex polygon with $O(n)$ vertices [Mat91]. By Theorem 4.2, $\mu(q) \geq 1 - 1/n^2$ for all points $q \in \mathcal{T}$, provided that the constant c is chosen appropriately. We also preprocess P for halfspace range reporting queries [CGL85]. \mathcal{T} can be computed in time $O(n \log^3 n)$ [Mat91], and constructing the half-plane range reporting data structure takes $O(n \log n)$ time [CGL85]. So the total preprocessing time is $O(n \log^3 n)$, and the size of the data structure is linear.

A query is answered as follows. Given a query point $q \in \mathbb{R}^2$, we first test in $O(\log n)$ time whether $q \in \mathcal{T}$. If the answer is yes, we simply return 1 as $\mu(q)$. If not, we compute in $O(\log n)$ time the two tangents ℓ_1, ℓ_2 of \mathcal{T} from q . For $i = 1, 2$, let $\xi_i = \ell_i \cap \mathcal{T}$, and let ℓ_i^- be the half-plane bounded by ℓ_i that does not contain \mathcal{T} . Set $\mathcal{P}_q = \mathcal{P} \cap (\ell_1^- \cup \ell_2^-)$ and $n_q = |\mathcal{P}_q|$. Let R_q be the subset of \mathcal{P}_q by choosing each point with probability γ .



By querying the half-plane range reporting data structure with each of these two tangent lines, we compute the set \mathcal{P}_q in time $O(\log n + n_q)$. Let $\omega_q = \Pr[q \notin \text{CH}(R_q \cup \mathcal{T})]$. We compute ω_q , in $(n_q \log n_q)$ time, by adapting the algorithm for computing $\mu(q)$ described in Section 2.

The correctness and efficiency of the algorithm follow from the following lemma, whose proof is omitted from this version.

Lemma 4.4. For any point $q \notin \mathcal{T}$, (i) $|\Pr[q \in \text{CH}(R_q \cup \mathcal{T})] - \mu(q)| \leq 1/n$; (ii) $n_q \leq 4t_0 = O(\gamma^{-1} \log n)$.

By Lemma 4.4, $n_q = O(\gamma^{-1} \log n)$, so the query takes $O(\gamma^{-1} \log(n) \log \log n)$ time. We thus obtain the following.

Theorem 4.5. Let \mathcal{P} be a set of n uncertain points in \mathbb{R}^2 in the unipoint model, where each point appears with probability γ . \mathcal{P} can be preprocessed in $O(n \log^3 n)$ time into a linear-size data structure so that for any point $q \in \mathbb{R}^2$, returns a value $\tilde{\mu}(q)$ in $O(\gamma^{-1} \log(n) \log \log n)$ time such that $|\tilde{\mu}(q) - \mu(q)| \leq 1/n$.

5. β -Hull

In this section, we consider the multipoint model, i.e., \mathcal{P} is a set of m uncertain point defined by the pairs $\{(P_1, \Gamma_1), \dots, (P_m, \Gamma_m)\}$. A convex set $C \subseteq \mathbb{R}^2$ is called β -dense with respect to \mathcal{P} if it contains β -fraction of each (P_i, Γ_i) , i.e., $\sum_{p_i^j \in C} \gamma_i^j \geq \beta$ for all $i \leq m$. The β -hull of \mathcal{P} , denoted by $\text{CH}_\beta(\mathcal{P})$, is the intersection of all convex β -dense sets with respect to \mathcal{P} . Note that for $m = 1$, $\text{CH}_\beta(\mathcal{P})$ is the set of points whose Tukey depth is at least $1 - \beta$. We first prove an $O(n)$ upper bound on the complexity of $\text{CH}_\beta(\mathcal{P})$ and then describe an algorithm for computing it.

Theorem 5.1. Let $\mathcal{P} = \{(P_1, \Gamma_1), \dots, (P_m, \Gamma_m)\}$ be a set of m uncertain points in \mathbb{R}^2 under the multipoint model with $P = \bigcup_{i=1}^m P_i$ and $|P| = n$. For any $\beta \in [0, 1]$, $\text{CH}_\beta(\mathcal{P})$ has $O(n)$ vertices.

Proof: We call a convex β -dense set C *minimal* if there is no convex β -dense set C' such that $C' \subset C$. A minimal convex β -dense set C is the convex hull of $P \cap C$. Therefore C is a convex polygon whose vertices are a subset of P . Obviously $\text{CH}_\beta(\mathcal{P})$ is the intersection of minimal convex β -dense sets. Therefore each edge of $\text{CH}_\beta(\mathcal{P})$ lies on a line passing through a pair of points of P , i.e., $\text{CH}_\beta(\mathcal{P})$ is the intersection of a set H of halfplanes whose bounding line passes through a pair of points of P . Next we argue that $|H| \leq 2n$.

Fix a point $p \in P$. We claim that H contains at most two halfplanes whose bounding lines pass through p . Indeed if $p \in \text{int}(\text{CH}_\beta(\mathcal{P}))$, then no bounding line of H passes through p ; if $p \in \partial(\text{CH}_\beta(\mathcal{P}))$, then at most two bounding lines of H pass through p ; and if $p \notin \text{CH}_\beta(\mathcal{P})$, then there are two tangent to $\text{CH}_\beta(\mathcal{P})$ from p . Hence at most two bounding lines of H pass through p , as claimed. ■

We describe a property of the set of lines supporting the edges of $\text{CH}_\beta(\mathcal{P})$, which will be useful for computing $\text{CH}_\beta(\mathcal{P})$. We call a line ℓ passing through a point $p \in P_i$ β -tangent of P_i at p if one of the open half-planes bounded by ℓ contains less than β -fraction of points of P_i but the corresponding closed half-plane contains at least β -fraction of points. Using a simple perturbation argument, the following can be proved.

Lemma 5.2. *A line supporting an edge of $\text{CH}_\beta(\mathcal{P})$ is β -tangent at two points of P .*

Algorithm. We describe the algorithm for computing the upper boundary \mathcal{U} of $\text{CH}_\beta(\mathcal{P})$. The lower boundary of $\text{CH}_\beta(\mathcal{P})$ can be computed analogously. It will be easier to compute \mathcal{U} in the dual plane. Let \mathcal{U}^* denote the dual of \mathcal{U} .

Recall that the dual of a point $p = (a, b)$ is the line $p^* : y = ax - b$, and the dual of a line $\ell : y = mx + c$ is the point $\ell^* = (m, -c)$. The point p lies above/below/on the line ℓ if and only if the dual point ℓ^* lies above/below/on the dual line p^* . Set $P_i^* = \{p_i^{j*} \mid p_i^j \in P_i\}$ and $P^* = \bigcup_{i=1}^m P_i^*$. For a point $q \in \mathbb{R}^2$ and for $i \leq m$, let $\kappa(q, i) = \sum \gamma_i^j$ where the summation is taken over all points $p_i^j \in P_i$ such that q lies below the dual line p_i^{j*} . We define the β -level Λ_i of P_i^* to be the upper boundary of the region $\{q \in \mathbb{R}^2 \mid \kappa(q, i) \geq \beta\}$. Λ_i is an x -monotone polygonal chain composed of the edges of the arrangement $\mathcal{A}(P_i^*)$. Further, the dual line of a point on Λ_i is a β -tangent line of P_i . Let Λ be the lower envelope of $\Lambda_1, \dots, \Lambda_m$.

Let ℓ be the line supporting an edge of \mathcal{U} . Using Lemma 5.2, it can be argued that the dual point ℓ^* is a vertex of Λ . Next, let q be a vertex of \mathcal{U} , then q cannot lie above any β -tangent line of any P_i , which implies that the dual line q^* passes through a pair of vertices of Λ and does not lie below any vertex of Λ . Hence, each vertex of \mathcal{U} corresponds to an edge of the upper boundary of the convex hull of Λ . By Theorem 5.1, \mathcal{U}^* , the dual of \mathcal{U} , has $O(n)$ vertices.

We now describe the algorithm for computing \mathcal{U}^* , which is similar to the one used for computing the convex hull of a level in an arrangement of lines [ASW08, Mat91]. We begin by describing a simpler procedure, which will be used as a subroutine in the overall algorithm.

Lemma 5.3. *Given a line ℓ , the intersection points of ℓ and Λ can be computed in $O(n \log n)$ time.*

Proof: We sort the intersections of the lines of P^* with ℓ . Let $\langle q_1, \dots, q_u \rangle$, $u \leq n$, be the sequence of these intersection points. For every $i \leq m$, $\kappa(q_1, i)$ can be computed in a total of $O(n)$ time. Given $\{\kappa(q_{j-1}, i) \mid 1 \leq i \leq m\}$, $\{\kappa(q_j, i) \mid 1 \leq i \leq m\}$ can be computed in $O(1)$ time. A point $q_j \in \Lambda$ if $q_j \in \Lambda_i$ for some i and lies below $\Lambda_{i'}$ for all other i' . This completes the proof of the lemma. ■

The following two procedures can be developed by plugging Lemma 5.3 into the parametric-search technique [ASW08, Mat91].

- (A) Given a point q , determine whether q lies above \mathcal{U}^* or return the tangent lines of \mathcal{U}^* from q . This can be done in $O(n \log^2 n)$ time.
- (B) Given a line ℓ , compute the edges of \mathcal{U}^* that intersect ℓ , in $O(n \log^3 n)$ time. (Procedure (B) uses (A) and parametric search.)

Given (B), we can now compute \mathcal{U}^* as follows. We fix a parameter $r > 1$ and compute a $(1/r)$ -cutting⁴ $\Xi = \{\Delta_1, \dots, \Delta_u\}$, where $u = O(r^2)$. For each Δ_i , we do the following. Using (B) we compute the edges of \mathcal{U}^* that intersect $\partial\Delta_i$. We can then deduce whether Δ_i contains any vertex of \mathcal{U}^* . If the answer is yes, we solve the problem recursively in Δ_i with the subset of lines of P^* that cross Δ_i . We omit the details from here and conclude the following.

Theorem 5.4. *Given a set \mathcal{P} of uncertain points in \mathbb{R}^2 under the multipoint model with a total of n sites, and a parameter $\beta \in [0, 1]$, the β -hull of \mathcal{P} can be computed in $O(n \log^3 n)$ time.*

Acknowledgments. P. Agarwal and W. Zhang are supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, and by an ERDC contract W9132V-11-C-0003. S. Har-Peled is supported by NSF grants CCF-09-15984 and CCF-12-17462. S. Suri and H. Yildiz are supported by NSF grants CCF-1161495 and CNS-1035917.

References

- [AAH⁺13] P. K. Agarwal, B. Aronov, S. Har-Peled, J. M. Phillips, K. Yi, and W. Zhang. Nearest neighbor searching under uncertainty II. In *Proc. 32nd ACM Sympos. Principles Database Syst.*, pages 115–126, 2013.
- [ACTY09] P. K. Agarwal, S.-W. Cheng, Y. Tao, and K. Yi. Indexing uncertain data. In *Proc. 28th ACM Sympos. Principles Database Syst.*, pages 137–146, 2009.
- [ASW08] P.K. Agarwal, M. Sharir, and E. Welzl. Algorithms for center and tverberg points. *ACM Trans. Algo.*, 5(1):5:1–5:20, December 2008.
- [CGL85] B. Chazelle, L. J. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25(1):76–90, 1985.
- [Cha93] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(1):145–158, 1993.
- [DRS09] N. N. Dalvi, C. Ré, and D. Suciu. Probabilistic databases: Diamonds in the dirt. *Commun. ACM*, 52(7):86–94, 2009.
- [Eck93] J. Eckhoff. Helly, radon, and Carathéodory type theorems. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, pages 389–448. North-Holland, 1993.
- [JLP11] A. Jørgensen, M. Löffler, and J. Phillips. Geometric computations on indecisive points. In *Proceedings of the 12th International Conference on Algorithms and Data Structures*, Proc. 12th Workshop Algorithms Data Struct., pages 536–547. Springer-Verlag, 2011.
- [KCS11a] P. Kamousi, T. M. Chan, and S. Suri. Closest pair and the post office problem for stochastic points. In *Proc. 12th Workshop Algorithms Data Struct.*, pages 548–559, 2011.

⁴A $(1/r)$ -cutting of P^* is a triangulation Ξ of \mathbb{R}^2 such that each triangle of Ξ crosses at most n/r lines of P^* .




Figure 3:

- [KCS11b] P. Kamousi, T.M. Chan, and S. Suri. Stochastic minimum spanning trees in euclidean spaces. In *Proc. 27th Annu. ACM Sympos. Comput. Geom.*, pages 65–74, 2011.
- [L09] M. Löffler. *Data Imprecision in Computational Geometry*. PhD thesis, Dept. Computer Sci., 2009.
- [Mat91] **J. Matoušek**. Computing the center of planar point sets. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Computational Geometry: Papers from the DIMACS Special Year*, pages 221–230. Amer. Math. Soc., 1991.
- [MS92] **J. Matoušek** and **O. Schwarzkopf**. Linear optimization queries. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 16–25, 1992.
- [Phi09] J.M. Phillips. *Small and Stable Descriptors of Distributions for Geometric Statistical Problems*. PhD thesis, Dept. Computer Sci., 2009.
- [Sei04] **R. Seidel**. Convex hull computations. In J. E. Goodman and **J. O’Rourke**, editors, *Handbook of Discrete and Computational Geometry*, pages 495–512. CRC Press, 2004.
- [SVY13] S. Suri, K. Verbeek, and H. Yildiz. On the most likely convex hull of uncertain points. In *Proc. 21st Annu. European Sympos. Algorithms*, pages 791–802, 2013.
- [ZYN12] Z. Zhao, D. Yan, and W. Ng. A probabilistic convex hull query tool. In *Proc. 15th Int. Conf. on Ext. Database Tech.*, pages 570–573, 2012.

A. Proof of Lemma 2.2

B. Computing Face Probabilities in Radial Order

Similar to the planar case, we can improve the computation time for the i th term to $O(n^{d-1})$ by considering the facets $f(P_\alpha)$ in radial order. In particular, let $L_\beta \subseteq P$ be a subset of $(d-2)$ sites. Let f_j denote the $(d-1)$ -dimensional simplex $f(L_\beta \cup \{q\} \cup \{p_j\})$ where $p_j \notin L_\beta$ and $p_j \neq p_i$. We can compute the probability that f_j is a facet of C for all facets f_j in constant amortized time as follows. We project all sites to the two-dimensional plane passing through q and orthogonal to the $(d-2)$ -dimensional hyperplane defined by $L_\beta \cup \{q\}$. (Such a plane is known as an orthogonal complement.) The hyperplane defined by $L_\beta \cup \{q\}$ projects onto q on this plane. Moreover, each facet f_j projects to a line segment extending from q . When we need to compute the probability that f_j is a facet of C , the set G_1 includes the sites on the other side of the line supporting f_j ’s projection with respect to p_i . (See Figure 3.) We compute probabilities for the facets f_j based on their radial order around q . The probability for the

next facet in the sweep can be computed by modifying the probability of the previous facet in constant amortized time as we have done for the planar case, as we can efficiently track how G_1 changes. As a final note, we point out that the total cost of all sorting involved is $O(n^{d-1} \log n)$ which is less than the overall cost of $O(n^d)$.

C. Membership Probability Algorithms in the Multipoint Model

C.1. The Planar Case

Let \mathcal{P} be an uncertain set of points in the multipoint model defined by site groups $\{P_1, \dots, P_m\}$. We denote the j th site in p_i by p_i^j and its probability by γ_i^j . For simplicity, we set $n_i = |P_i|$. We define P to be the set all sites, i.e. $P = \bigcup_{1 \leq i \leq m} p_i$, and set $n = |P| = \sum_{1 \leq i \leq m} n_i$. Under this setting, we want to compute the membership probability of a given point q . Recall that the sites from a single site group p_i are dependent, i.e., they cannot co-exist in an outcome $B \subseteq P$ of the probabilistic experiment.

The algorithm for the unipoint model easily extends to the multipoint model. The main difference is the way we compute the probability for an edge. The rest of the algorithm remains the mostly the same. We now see this in more depth.

Let V and C be defined as before. As in the unipoint model, q is in the convex hull of B if and only if $q \in V$. We follow a similar strategy and decompose $\Pr[q \in V]$ as follows:

$$\Pr[q \in V] = \Pr[B = \emptyset] + \sum_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n_i}} \Pr[qp_i^j \text{ is the witness edge of } q \notin \text{CH}(B)].$$

The first term is trivial to compute in $O(n)$ time. We compute the probability that qp_i^j forms a witness edge of B as follows. Let $G_{i,j}$ be the set of sites to the right of the line $\overleftrightarrow{qp_i^j}$ where the right direction is with respect to the vector $\overrightarrow{qp_i^j}$. As in the unipoint model, the segment qp_i^j is the witness edge of B if and only if $p_i^j \in B$ and $B \cap G_{i,j} = \emptyset$. We can write the corresponding probability as follows:

$$\begin{aligned} \Pr[p_i^j \in B \wedge B \cap G_{i,j} = \emptyset] &= \Pr[p_i^j \in B] \times \Pr[B \cap G_{i,j} = \emptyset \mid p_i^j \in B] \\ &= \Pr[p_i^j \in B] \times \prod_{1 \leq k \leq m} \Pr[B \cap G_{i,j} \cap P_k = \emptyset \mid p_i^j \in B] \\ &= \Pr[p_i^j \in B] \times \prod_{\substack{1 \leq k \leq m \\ k \neq i}} \Pr[B \cap P_k \cap G_{i,j} = \emptyset] \\ &= \gamma_i^j \times \prod_{\substack{1 \leq k \leq m \\ k \neq i}} \left(1 - \sum_{l \mid p_k^l \in G_{i,j}} \gamma_k^l \right). \end{aligned}$$

This expression can be easily computed in $O(n)$ time. It follows that one can compute $\overline{\mu(q)}$, thus $\mu(q)$, in $O(n^2)$ time.

As before, the computation time can be improved to $O(n \log n)$ by computing the witness edge probabilities in radial order around q . Let the circular sequence p'_1, p'_2, \dots, p'_n be the counter-clockwise order of all sites around q , where each p'_i is a distinct site p_a^b . We first compute the probability that qp'_1 is the witness edge in $O(n)$ time and also remember the values of the intermediate factors used in the computation. (The factors inside the $\prod_{1 \leq k \leq m}$ expression.) Then, for increasing values of i

from 2 to n , we compute the probability that $qp'(k)$ is the witness edge by updating the probability for $qp'(k-1)$. As a first step to this update, we update the values of the intermediate factors. To be more specific, let W_i denote the set of sites in the open wedge bounded by the lines $\overleftrightarrow{qp'(i)}$ and $\overleftrightarrow{qp'(i-1)}$. Also, for simplicity, assume that $p'(k) = p_a^b$ and $p'(k-1) = p_c^d$. Notice that $G_{a,b} = G_{c,d} \cup \{p_c^d\} \setminus W_i$. Then, for each site p_k^l in W_i , the k th factor increases by γ_k^l . Also, the c th factor decreases by γ_c^d . Finally, we temporarily set the value of the a th factor to 1 (to cover the case $k \neq i$ in the expression). Then, we can compute the witness edge probability for $qp'(k)$ by multiplying the probability of $qp'(k-1)$ with γ_a^b/γ_c^d and the multiplicative change in each intermediate factor. The cost of a single update is $O(1)$ amortized, as each site can contribute to at most 4 updates as in the unipoint case.

C.2. The d -dimensional case

All of the arguments in the d -dimensional algorithm are also easily extended to the multipoint. As before, we compute $\mu(q)$ by computing the probability $\Pr[q \in V]$. Following the same strategy, we decompose it as

$$\begin{aligned} \Pr[q \in V] &= \Pr[q = \lambda(B \cup \{q\})] \\ &+ \sum_{1 \leq i \leq m} \left(\sum_{1 \leq j \leq n_i} \Pr[p_i^j = \lambda(B \cup \{q\}) \wedge q \in V] \right). \end{aligned}$$

It is trivial to compute the first term in $O(n)$ time. We now show how to compute each term inside the summations in $O(n^{d-1})$ time. This implies a total time of $O(n^d)$.

Clearly, Lemma 2.2 extends to the multipoint model, so we can use p_i^j -escaping facets to decompose our probability. Given a subset of sites $P_\alpha \subseteq P \setminus \{p_i^j\}$ of size $(d-1)$, define $f(P_\alpha)$ to be the $(d-1)$ -dimensional simplex whose vertices are the points in P_α and q . Then,

$$\begin{aligned} \Pr[p_i^j = \lambda(B \cup \{q\}) \wedge q \in V] &= \Pr[p_i^j = \lambda(B \cup \{q\}) \wedge q' \in V'] \\ &+ \sum_{\substack{P_\alpha \subseteq P \setminus \{p_i^j\} \\ |P_\alpha| = (d-1) \\ f(P_\alpha) \text{ is } p_i^j\text{-escaping for } q}} \Pr[p_i^j = \lambda(B \cup \{q\}) \wedge f(P_\alpha) \text{ is a facet of } C]. \end{aligned}$$

The first term is computed recursively. We compute each term of the summation as follows. Let I_α be the set of group indices of the sites in P_α , i.e., $I_\alpha = \{u \mid \exists v. p_u^v \in P_\alpha\}$. As before, let $G_1 \subseteq P$ be the subset of sites which are on the other side of the hyperplane supporting $f(P_\alpha)$ with respect to p_i^j . Let $G_2 \subseteq P$ be the subset of sites that are below p_i^j along the x_d -axis. Following the same strategy, we write the desired probability as the probability that all points in P_α and p_i^j exist in B , and all points in $G_1 \cup G_2$ are absent from B . This probability is clearly zero, if any of the following conditions hold:

- $P_\alpha \cap G_2 \neq \emptyset$.
- p_i^j has a higher x_d -coordinate than q .
- P_α contains any two sites from the same uncertain point P_k .
- P_α contains any site from p_i .

Otherwise, we can write the probability as follows:

$$\begin{aligned}
& \Pr[p_i^j \in B \wedge P_\alpha \cap B = P_\alpha \wedge B \cap (G_1 \cup G_2) = \emptyset] \\
&= \Pr[p_i^j \in B] \times \Pr[P_\alpha \cap B = P_\alpha \mid p_i^j \in B] \times \\
&\quad \Pr[B \cap (G_1 \cup G_2) = \emptyset \mid p_i^j \in B \wedge P_\alpha \cap B = P_\alpha] \\
&= \Pr[p_i^j \in B] \times \Pr[P_\alpha \cap B = P_\alpha] \times \\
&\quad \Pr[B \cap (G_1 \cup G_2) = \emptyset \mid p_i^j \in B \wedge P_\alpha \cap B = P_\alpha] \\
&= \Pr[p_i^j \in B] \times \Pr[P_\alpha \cap B = P_\alpha] \times \\
&\quad \prod_{\substack{1 \leq u \leq m \\ u \neq i \\ u \notin I_\alpha}} \left(\Pr[P_u \cap B \cap (G_1 \cup G_2) = \emptyset] \right) \\
&= \gamma_i^j \times \prod_{u, v \mid p_u^v \in P_\alpha} \gamma_u^v \times \prod_{\substack{1 \leq u \leq m \\ u \neq i \\ u \notin I_\alpha}} \left(1 - \sum_{v \mid p_u^v \in (G_1 \cup G_2)} \gamma_u^v \right).
\end{aligned}$$

The expression takes linear time to compute and thus summation term can be computed in $O(n^d)$ time. Then, by induction, the computation of the term for the site p_i^j takes $O(n^d)$ time. As before, we can improve the computation time each term to $O(n^{d-1})$ by considering the facets $f(P_\alpha)$ in radial order. This implies a total complexity of $O(n^d)$ for the algorithm.

D. Computing the Probability Map in $O(n^4)$ Time

In this section, we describe how to compute the **probability map** for a given set of uncertain points on the plane in $O(n^4)$ time, rather than $O(n^5 \log n)$. For simplicity, we show how to compute the probability associated with each **face** (or cell) of the probability map. Our algorithm can easily be adapted to compute the probabilities of edges and vertices as well, by taking care of degeneracies. Also, we assume that the input is given in the unipoint model, however, we briefly explain how to extend the algorithm to the multipoint model.

The high level idea of our algorithm is as follows. Recall that the structure of the probability map is an arrangement of $O(n^2)$ lines, containing $O(n^4)$ faces. We first compute the membership probability of one of the faces, say F , in $O(n \log n)$ time. We then compute the membership probabilities of the faces neighboring F , in $O(1)$ time per each, by modifying the probability of F . By repeatedly applying the same idea of expanding into the neighbors, we can compute the probability of all faces in $O(n^4)$ time.

To complete our algorithm, we now describe how we can compute the probability of a face F' by using the already computed probability of one of its neighbors F . Without loss of generality, assume that F and F' are separated by a vertical line passing through the sites p_i and p_j and F is to the left of F' . Notice that the boundary separating F and F' is only a segment of the vertical line and does not contain p_i or p_j . Now imagine that a point q moves through this boundary, crossing from F to F' . It is easy to see that the change in the membership probability of q is due to the changes in witness edge probabilities of the segments qp_i and qp_j , as other sites are irrelevant. We now describe the change in the witness edge probability of qp_i . The probability of qp_j changes analogously. The change in the probability of qp_i happens differently for two cases:

1. p_i **is above** p_j : Then, p_j switches from the right side of the line $\overleftrightarrow{qp_i}$ to its left side (where right direction is with respect to the vector $q\vec{p_i}$). Consequently, the probability of qp_i changes by a factor of $\frac{1}{\gamma_j}$.
2. p_i **is below** p_j : Then, p_j switches from the left side of the line $\overleftrightarrow{qp_i}$ to its right side. Consequently, the probability of qp_i changes by a factor of $\overline{\gamma_j}$.

The changes clearly require constant time operations, and thus the membership probability of F' can be computed in $O(1)$ time.

The extension of this technique to the multipoint model is straightforward. The only major difference is that we need to remember (similar to what is done in Appendix C) the intermediate factors when switching from face to face, as updating the witness edge probabilities requires updating these factors first. The total cost of an update remains $O(1)$ because each face switch updates one intermediate factor of two witness edge probabilities.